

Deep Two-View Structure-from-Motion Revisited

*Jianyuan Wang¹, *Yiran Zhong¹, Yuchao Dai², Stan Birchfield³,
Kaihao Zhang¹, Nikolai Smolyanskiy³, Hongdong Li¹

¹Australian National University, ²Northwestern Polytechnical University, ³NVIDIA

Abstract

Two-view structure-from-motion (SfM) is the cornerstone of 3D reconstruction and visual SLAM. Existing deep learning-based approaches formulate the problem by either recovering absolute pose scales from two consecutive frames or predicting a depth map from a single image, both of which are ill-posed problems. In contrast, we propose to revisit the problem of deep two-view SfM by leveraging the well-posedness of the classic pipeline. Our method consists of 1) an optical flow estimation network that predicts dense correspondences between two frames; 2) a normalized pose estimation module that computes relative camera poses from the 2D optical flow correspondences, and 3) a scale-invariant depth estimation network that leverages epipolar geometry to reduce the search space, refine the dense correspondences, and estimate relative depth maps. Extensive experiments show that our method outperforms all state-of-the-art two-view SfM methods by a clear margin on KITTI depth, KITTI VO, MVS, Scenes11, and SUN3D datasets in both relative pose and depth estimation.

1. Introduction

Two-view structure-from-motion (SfM) is the problem of estimating the camera motion and scene geometry from two image frames of a monocular sequence. As the foundation of both 3D reconstruction and visual simultaneous localization and mapping (vSLAM), this important problem finds its way into a wide range of applications, including autonomous driving, augmented/virtual reality, and robotics.

Classic approaches to two-view SfM follow a standard pipeline of first matching features/edges between the two images, then inferring motion and geometry from those matches [8, 22, 30, 14, 12, 11, 47, 52]. When imaging conditions are well-behaved (constant lighting, diffuse and rigid surfaces, and non-repeating visual texture), the match-

ing process is well-posed. And, once the matches have been found, the motion and geometry can be recovered.

For decades, researchers who work in this area have generally required at least two views, and their methods have recovered only *relative* camera motion and *relative* scene geometry (that is, shape up to an unknown scale factor). Without a *priori* knowledge of scale or recognizable objects in the scene, it is impossible to recover scene geometry from a single view [18, 48, 44, 25]. Similarly, it is impossible to infer *absolute* scale from two views of a scene [17].

With the rise of deep learning, a number of researchers have recently explored neural network-based solutions to two-view SfM. Most of these methods fall into one of two categories. In the first category, which we shall call Type **I**, the problem is treated as a joint optimization task of monocular depth and pose regression [53, 46, 32]. Two networks are used: one to estimate the *up-to-scale* depth from a single image, and the other one to predict the *up-to-scale* camera pose from two input images. Both networks act independently during inference. In the second category, denoted Type **II**, the *scaled* camera pose and the *scaled* depth are inferred from the image pair, and are iteratively refined via multi-view geometry [39, 35, 36]. While the power of deep learning allows both Type **I** and Type **II** solutions to achieve compelling results, we note that their formulations attempt to solve one of the *ill-posed* [3] problems mentioned above.

In this paper, we revisit the use of deep learning for two-view SfM. Our framework follows the classic SfM pipeline that features are matched between image frames to yield *relative* camera poses, from which *relative* depths are then estimated. By combining the strengths of deep learning within a classic pipeline, we are able to avoid ill-posedness, which allows our approach to achieve state-of-the-art results on several benchmarks.

A comparison between our approach and existing pipelines is shown in Fig. 1. Our method operates by first estimating dense matching points between two frames using a deep optical flow network [50, 40], from which a set of highly reliable matches are sampled in order to compute the relative camera pose via a GPU-accelerated classic five-point algorithm [24] with RANSAC [13]. Since these

* indicates equal contribution, listed in alphabetical order. Yiran is the corresponding author. Work was partially done when Yiran was an intern at NVIDIA, Redmond, WA.

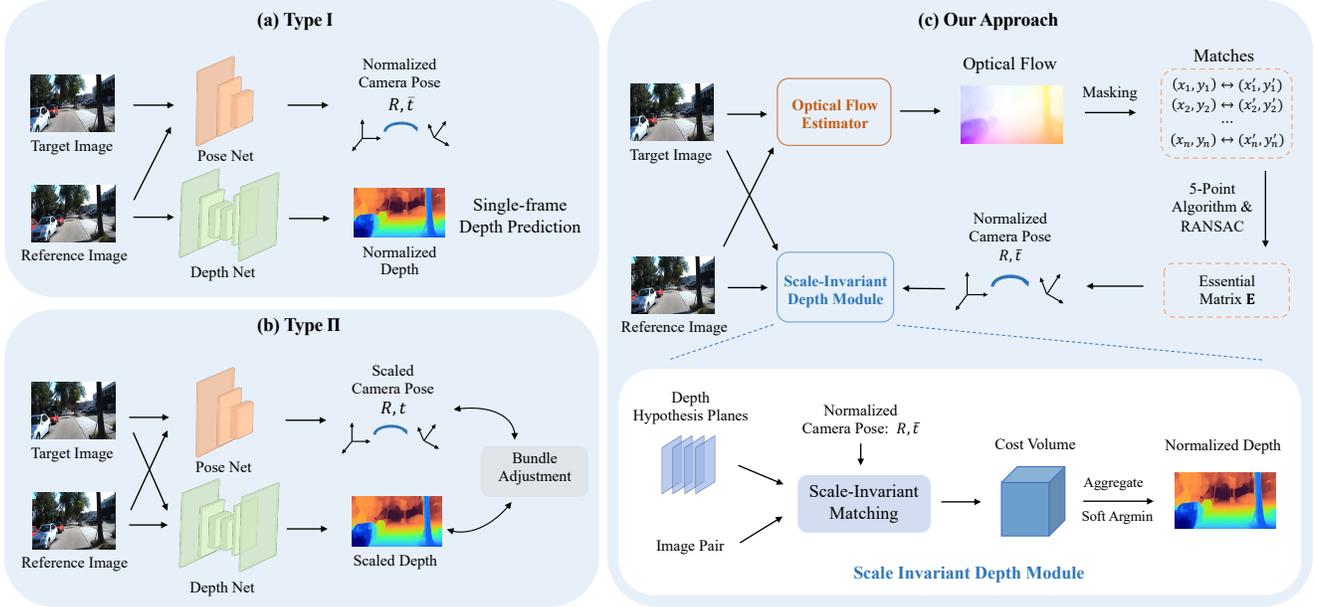


Figure 1. **Comparison between our method and previous deep monocular structure-from-motion methods.** We formulate camera pose estimation as a 2D matching problem (optical flow) and depth prediction as a 1D matching problem along an epipolar line. In contrast, previous methods suffer from ill-posedness (either single-frame depth prediction, in the case of Type I, or scaled estimates, in the case of Type II).

relative camera poses have scale ambiguity, the estimated depth suffers from scale ambiguity as well. Therefore, in order to supervise the estimated scale-ambiguous depth with the (scaled) ground truth depth, we propose a scale-invariant depth estimation network combined with scale-specific losses to estimate the final relative depth maps. Since the search space of the depth estimation network is reduced to epipolar lines thanks to the camera poses, it yields higher accuracy than directly triangulating the optical flows with the estimated camera poses. We demonstrate the effectiveness of our framework by achieving state-of-the-art accuracy in both pose and depth estimation on KITTI depth, KITTI VO, MVS, Scenes11, and SUN3D datasets.

Our main contributions are summarized as:

- 1) We revisit the use of deep learning in SfM, and propose a new deep two-view SfM framework that avoids ill-posedness. Our framework combines the best of deep learning and classical geometry.
- 2) We propose a scale-invariant depth estimation module to handle the mismatched scales between ground truth depth and the estimated depth.
- 3) Our method outperforms all previous methods on various benchmarks for both relative pose estimation and depth estimation under the two-view SfM setting.

2. Two-view Geometry: Review and Analysis

The task of two-view SfM refers to estimating the relative camera poses and dense depth maps from two consecutive monocular frames. In classic geometric vision, it is

well understood that the camera poses as well as the depth maps can be computed from image matching points alone without any other information [26].¹

Given a set of image matching points in homogeneous coordinates, $\mathbf{x}_i = [x_i \ y_i \ 1]^T$ and $\mathbf{x}'_i = [x'_i \ y'_i \ 1]^T$ with known camera intrinsic matrix \mathbf{K} , the two-view SfM task is to find a camera rotation matrix \mathbf{R} and a translation vector \mathbf{t} as well as the corresponding 3D homogeneous point \mathbf{X}_i such that:

$$\mathbf{x}_i = \mathbf{K} \begin{bmatrix} \mathbf{I} & \mathbf{0} \end{bmatrix} \mathbf{X}_i \quad \mathbf{x}'_i = \mathbf{K} \begin{bmatrix} \mathbf{R} & \mathbf{t} \end{bmatrix} \mathbf{X}_i \quad \forall i. \quad (1)$$

A classical method to solve this problem consists of three consecutive steps: 1) Computing the essential matrix \mathbf{E} from the image matching points \mathbf{x}_i and \mathbf{x}'_i ; 2) Extracting the relative camera pose \mathbf{R} and \mathbf{t} from the essential matrix \mathbf{E} ; 3) Triangulating the matching points \mathbf{x}_i and \mathbf{x}'_i with the camera pose to get the 3D point \mathbf{X}_i .

All steps in this pipeline are *well-posed* problems. The essential matrix \mathbf{E} can be solved with at least 5 matching points using the equation below:

$$\mathbf{x}'_i{}^T \mathbf{K}^{-T} \mathbf{E} \mathbf{K}^{-1} \mathbf{x}_i = 0 \quad \forall i. \quad (2)$$

\mathbf{R} and \mathbf{t} can be computed from \mathbf{E} using matrix decomposition such that $\mathbf{E} = \mathbf{S}\mathbf{R}$, where \mathbf{S} is a skew symmetric matrix and \mathbf{R} is a rotation matrix. Since for any non-zero scaling factor α , $[\alpha\mathbf{t}]_{\times} \mathbf{R} = \alpha [\mathbf{t}]_{\times} \mathbf{R} = \alpha \mathbf{E}$ provides a

¹Excluding degenerate cases.

valid solution, there is a scale ambiguity for relative camera pose estimation. The 3D point X_i can be computed by triangulation with a global scale ambiguity.

The method above assumes the ideal case in which all image points are perfectly matched. To handle mismatched points in real scenarios, researchers have established a classical standard pipeline to estimate geometry information from two consecutive frames [17].

2.1. The Classic Standard Pipeline

With decades of development and refinement, the classic standard pipeline [17] is widely used in many conventional state-of-the-art SfM and vSLAM systems [30, 33, 1]. Since almost all geometry information can be recovered from image matching points, the key is to recover a set of (sparse or dense) accurate matching points. To this end, the pipeline often starts with sparse (or semi-dense) distinct feature extraction and matching to get sparse matching points, as sparse matching is more accurate than dense matching. To further refine the matching results, the RANSAC scheme [13] is used to filter the matching points that do not fit the majority motion. These outliers often include mismatches and dynamic objects in a scene. After retrieving the camera poses from the refined matching points, the depth of these points can be computed via triangulation. In some cases, if it is desired to estimate dense depth maps rather than the sparse 3D points, multi-view stereo matching algorithms can be used to recover the dense depth maps with the estimated camera poses.

The Achilles' heel of this pipeline is therefore the matching of points. Conventional matching algorithms often suffer from low accuracy on non-Lambertian, blurry, and textureless surfaces. However, this shortage can be largely alleviated by deep learning [34, 37, 40, 51, 6, 23]. With sufficient training data, such networks can learn to handle these scenarios. In our proposed approach, we leverage a deep optical flow network [40] to compute these correspondences.

2.2. Deep Learning based Methods

As discussed earlier, two-view SfM requires to estimate both camera poses and dense depth maps. Existing deep learning based methods either formulate the problem as pose and monocular depth regression (Type I) or as pose regression and multi-view stereo matching (Type II). We analyze both types of methods below.

Type I methods consist of a monocular depth estimation network and a pose regression network. The two-view geometry constraints are used as self-supervisory signals to regularize both camera poses and depth maps [53, 46, 28, 5, 2, 32]. As a result, most of these approaches are self-supervised. Because single-view depth estimation is inherently ill-posed, as discussed earlier, these methods are fun-

damentally limited by how well they can solve that challenging problem. They rely on the priors from the training data to predict depth only given a single image.

Moreover, since the two-view geometry constraints are only suitable for a stationary scene, SfMLearner [53] simultaneously estimates an explainability mask to exclude the dynamic objects while GeoNet [46] utilizes an optical flow module to mask out these outliers by comparing the rigid flow (computed by camera poses and depth maps) with the non-rigid flow (computed by the optical flow module). Other methods focus on implementing more robust loss functions, such as ICP loss [28], motion segmentation loss [32], or epipolar loss [5].

Type II methods require two image frames to estimate depth maps and camera poses at test time (unlike Type I methods, which estimate depth from a single frame). Most supervised deep methods fall into this category. As a pioneer of this type, DeMoN [39] concatenates a pair of frames and uses multiple stacked encoder-decoder networks to regress camera poses and depth maps, implicitly utilizing multi-view geometry.

Similar strategies have been adapted by [35, 7, 41, 36] through replacing generic layers between camera poses and depth maps with optimization layers that explicitly enforce multi-view geometry constraints. For example, BANet [35] parameterizes dense depth maps with a set of depth bases [49] and imposes bundle adjustment as a differentiable layer into the network architecture. Wang *et al.* [41] use regressed camera poses to constrain the search space of optical flow, estimating dense depth maps via triangulation. DeepV2D [36] separates the camera pose and depth estimation, iteratively updating them by minimizing geometric reprojection errors. Similarly, DeepSfM [42] initiates its pose estimation from DeMoN [39], sampling nearby pose hypotheses to bundle adjust both poses and depth estimation. Nevertheless, with the ground truth depth as supervision, it requires the pose regression module to estimate camera poses with absolute scale, which is generally impossible from a pair or a sequence of monocular frames alone [17]. To mitigate this ill-posed problem, they utilize dataset priors and semantic knowledge of the scene to estimate the absolute scale.

3. Method

In this section, we propose a new deep two-view SfM framework that aims to address the Achilles' heel of the classical SfM pipeline (*viz.*, matching) via deep learning. Our method is able to find better matching points and therefore more accurate poses and depth maps, especially for textureless and occluded areas. At the same time, it follows the wisdom of classic methods to avoid the ill-posed problems. By combining the best of both worlds, our approach is able to achieve state-of-the-art results, outperforming all previ-

ous methods by a clear margin.

Following the classic standard pipeline [17], we formulate the two-frame structure-from-motion problem as a three-step process: 1) match corresponding points between the frames, 2) estimate the essential matrix and hence the relative camera pose, and 3) estimate dense depth maps up to an unknown scale factor. These steps, along with the loss function used for training, are described in more detail in the following subsections.

3.1. Optical Flow Estimation

As a fundamental problem in computer vision, optical flow estimation has been extensively studied for several decades [20]. With the recent progress in deep learning, deep optical flow methods now dominate various benchmarks [16, 4] and can handle large displacements as well as textureless, occluded, and non-Lambertian surfaces. In our framework, we utilize the state-of-the-art network, DCL-Flow [40], to generate dense matching points between two consecutive frames. This method uses a displacement-invariant matching cost learning strategy and a soft-argmin projection layer to ensure that the network learns dense matching points rather than image-flow regression. The network was trained on synthetic datasets (FlyingChairs [9] and FlyingThings [29]) to avoid data leakage, *i.e.*, the network was not trained on any of the test datasets.

3.2. Essential Matrix Estimation

The traditional approach to estimating camera pose between two image frames is to match sparse points, *e.g.*, SIFT features [27]. Then, given a set of matching points $\mathbf{x} \leftrightarrow \mathbf{x}'$ and the camera intrinsic matrix \mathbf{K} , the essential matrix [26] \mathbf{E} can be recovered from the five-point algorithm [31, 24]. By decomposing the essential matrix as $\mathbf{E} = [\mathbf{t}]_{\times} \mathbf{R}$, the rotation matrix \mathbf{R} and the translation vector \mathbf{t} can be recovered up to a scale ambiguity. Conventionally, outliers in the matching points are filtered using robust fitting techniques such as RANSAC [13]. RANSAC repeatedly estimates the essential matrix from randomly sampled minimal matching sets and selects the solution that is satisfied by the largest proportion of matching points under a certain criterion.

Unlike all previous deep learning-based methods that regress the camera poses from input images, we use matching points to compute the camera poses. The key question is this: How to robustly filter the noisy dense matches from optical flow in order to retain only the high quality matches? There are multiple ways to filter out unreliable matching points such as flow uncertainty, consistency check, or using a network to regress a mask. Empirically, we find that simply using SIFT keypoint locations (note that we do *not* use SIFT matching) to generate a mask works well in all datasets. The hypothesis is that optical flow is more ac-

curate in rich textured areas. The optical flow matches at the locations within the mask are filtered by RANSAC with GPU acceleration, to avoid distraction by dynamic objects. After retrieving the essential matrix \mathbf{E} , the camera pose (\mathbf{R}, \mathbf{t}) is recovered using matrix decomposition.

3.3. Scale-Invariant Depth Estimation

Once we have recovered the up-to-scale relative camera pose, with the dense matching points from optical flow estimation, we could compute the dense depth map by performing triangulation. However, such an approach would not take advantage of the epipolar constraint. As a result, we perform the matching again by constraining the search space to epipolar lines computed from the relative camera poses. This process is similar to multi-view stereo (MVS) matching with one important difference: we do not have the absolute scale in inference. With the up-to-scale relative pose, if we were to directly supervise the depth estimation network with ground truth depth, there would be a mismatch between the scale of the camera motion and the scale of the depth map.

Previous approaches. To resolve this paradox, previous methods either use a scale-invariant loss [10] or regress the absolute scale with a deep network [39, 36, 42].

The scale-invariant loss ℓ_{SI} is defined as:

$$\ell_{\text{SI}} = \sum_{\mathbf{x}} \left(\log(d_{\mathbf{x}}) - \log(\hat{d}_{\mathbf{x}}) + \eta(d, \hat{d}) \right)^2, \quad (3)$$

where $d_{\mathbf{x}}$ and $\hat{d}_{\mathbf{x}}$ are the ground truth and estimated depth, respectively, at pixel \mathbf{x} ; and $\eta(d, \hat{d}) = \frac{1}{N} \sum_{\mathbf{x}} (\log(\hat{d}_{\mathbf{x}}) - \log(d_{\mathbf{x}}))$, where N is the number of pixels, measures the mean log difference between the two depth maps. While working for the direct depth regression pipelines, the scale-invariant loss introduces an ambiguity for network learning as the network could output depth maps with different scales for each sample. This loss may hinder the principle of plane-sweep, where depth maps with consistent scale across frames of the sequence are desired. Plane sweep [21] is the process that enforces the epipolar constraint, which reduces the search space from 2D to 1D.

Plane-sweep powered networks require consistent scale during the training and testing process. For example, if we train a network with absolute scales and test it with a normalized scale, its performance will drop significantly (we provide an ablation study in Section 4.4). Since it is impossible to recover the absolute scale from two images, some previous methods [39, 36, 42] use a network to regress a scale to mimic the absolute scale in inference. This strategy slightly alleviates the scale paradox at the cost of making the problem ill-posed again.

Scale-Invariant Matching. To solve this paradox and keep the problem well-posed, we propose a scale-invariant

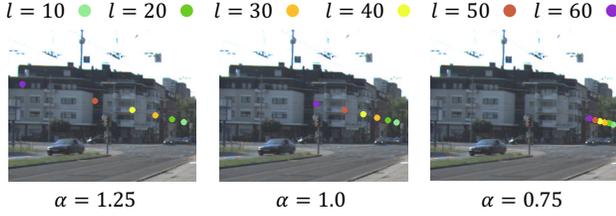


Figure 2. **The Effect of Various Scale Factors during Plane Sweep.** For a certain pixel, we visualize its six depth hypotheses with different colors in the target frame. As the scale factor α changes, the sampling distribution varies.

matching process to recover the up-to-scale dense depth map. Mathematically, given an image point \mathbf{x} , we generate L matching candidates $\{\mathbf{x}'_l\}_{l=1}^L$:

$$\mathbf{x}'_l \sim \mathbf{K}[\mathbf{R} \mid \mathbf{t}] \begin{bmatrix} (\mathbf{K}^{-1}\mathbf{x})d_l \\ 1 \end{bmatrix}, \quad (4)$$

where $d_l = (L \times d_{\min})/l$, ($l = 1, \dots, L$) is the depth hypothesis and d_{\min} is a fixed minimum depth.

In the standard plane-sweep setting, the sampling distribution of matching candidates varies depending on the scale factor $\alpha = \|\mathbf{t}\|_2$, as illustrated in Fig. 2. Instead, we normalize the translation vectors to $\bar{\mathbf{t}} \equiv \mathbf{t}/\alpha$ such that $\|\bar{\mathbf{t}}\|_2 = 1$, since we do not know the absolute scale in our problem. Substituting the normalized translation $\bar{\mathbf{t}}$ for \mathbf{t} in Eq. (4), with fixed $\{d_l\}_{l=1}^L$, the distribution of matching candidates $\{\mathbf{x}'_l\}_{l=1}^L$ are now invariant to scale.

To make the estimated and ground truth depths compatible, according to Eq. (4), we need to scale the estimated depth $\hat{\mathbf{d}}$ correspondingly to match the ground truth depth \mathbf{d} :

$$\mathbf{d} \sim \alpha_{\text{gt}} \hat{\mathbf{d}}, \quad (5)$$

where α_{gt} refers to the ground truth scale.

This scale-invariant matching strategy plays a crucial role in our framework as it makes our network no longer suffer from the scale misalignment problem. Please note our competitors cannot benefit from scale-invariant matching because they usually avoid the scale misalignment problem by predicting absolute scales. A detailed discussion is provided in Section 4.4.

3.4. Loss Function

Our framework is trained in an end-to-end manner with the supervision of ground truth depth maps and ground truth scales. Given a predicted depth $\hat{\mathbf{d}}$ and a ground truth depth \mathbf{d} , we supervise the depth using the Huber loss:

$$\mathcal{L}_{\text{depth}} = \sum_{\mathbf{x}} \ell_{\text{huber}}(\alpha_{\text{gt}} \hat{\mathbf{d}}_{\mathbf{x}} - \mathbf{d}_{\mathbf{x}}), \quad (6)$$

where $\ell_{\text{huber}}(z) = 0.5z^2$ if $|z| < 1$, $|z - 0.5|$ otherwise. It should be noted that our predicted depth is up-to-scale and does not require ground truth scale at inference time.

If both ground truth camera pose (\mathbf{R}, \mathbf{t}) and ground truth depth $d_{\mathbf{x}}$ are given, we can also update the optical flow network by computing the rigid flow $\mathbf{u}_{\mathbf{x}} \equiv \mathbf{x}' - \mathbf{x}$ for 2D point \mathbf{x} :

$$\mathbf{x}' \sim \mathbf{K}[\mathbf{R} \mid \mathbf{t}] \begin{bmatrix} (\mathbf{K}^{-1}\mathbf{x})d_{\mathbf{x}} \\ 1 \end{bmatrix}. \quad (7)$$

The rigid flow can work as a supervision signal, computing the ℓ_2 distance with the estimated optical flow $\hat{\mathbf{u}}_{\mathbf{x}}$:

$$\mathcal{L}_{\text{flow}} = \sum_{\mathbf{x}} (\hat{\mathbf{u}}_{\mathbf{x}} - \mathbf{u}_{\mathbf{x}})^2. \quad (8)$$

The total loss function of our framework is then given by:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{depth}} + \lambda \mathcal{L}_{\text{flow}}. \quad (9)$$

We set $\lambda = 1$ to fine-tune the optical flow estimator, or $\lambda = 0$ to use the flow model pretrained on synthetic datasets.

4. Experiments

In this section, we provide quantitative and qualitative results of our framework on various datasets, showing comparison with state-of-the-art SfM methods. We also provide an extensive ablation study to justify our framework design. Due to the scale-ambiguity nature of the two-view SfM problem, we scale the results of ours and others using the same scaling strategy as in [36, 35]. For all experiments, our optical flow estimator is [40] while the architecture of depth estimator is based on [21]. Implementation details (such as hyperparameters of the optimizer or network) are provided in the supplementary material.

4.1. Datasets

KITTI Depth [16] is primarily designed for monocular depth evaluation in autonomous driving scenarios, which does not take camera motions and dynamic objects into account. The Eigen split [10], which contains 697 single frames for testing, is a widely used split for evaluating monocular depth estimation. To adapt it for two-view SfM evaluation, we pair nearby frames. Also, since the Eigen split contains a number of frames with nearly static camera motion or many moving objects (which lead to ill-posed situations for two-view SfM), we filter out these frames to create an Eigen SfM split (256 frames) to evaluate SfM algorithms in well-conditioned scenarios. Specifically, we first pair each frame with its next frame then manually remove these pairs with small relative translations (less than 0.5 meters) or contain large dynamic objects².

KITTI VO [16] is primarily used for evaluating camera pose estimation. It contains ten sequences (more than 20k frames) with ground truth camera poses. According to the

²We define a dynamic object which occupies more than 20% pixels of a scene as a large dynamic object.

Table 1. **Depth Evaluation on KITTI Depth Dataset.** We compare our results to state-of-the-art single-frame depth estimation methods and deep SfM methods on the KITTI depth Eigen split. We evaluate all SfM methods under two-view SfM setting for a fair comparison. The ‘‘Eigen SfM’’ split (256 frames) excludes frames that are close to static or contain many dynamic objects in the Eigen split. The type **S** means supervised single frame depth estimation. Note that Type **I** methods are self-supervised methods. Bold indicates the best.

Split	Type	Method	lower is better					higher is better		
			Abs Rel	Sq Rel	RMSE	RMSE _{log}	D1-all	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Eigen	S	DORN [15]	0.072	0.307	2.727	0.120	0.163	0.932	0.984	0.994
		VNL [45]	0.072	-	3.258	0.117	0.176	0.938	0.990	0.998
	I	SfMLearner [53]	0.208	1.768	6.856	0.283	-	0.678	0.885	0.957
		GeoNet [46]	0.155	1.296	5.857	0.233	-	0.793	0.931	0.973
		CCNet [32]	0.140	1.070	5.326	0.217	-	0.826	0.941	0.975
		GLNet [5]	0.099	0.796	4.743	0.186	-	0.884	0.955	0.979
	II	BANet [35]	0.083	-	3.640	0.134	-	-	-	-
		DeepV2D [36]	0.064	0.350	2.946	0.120	0.142	0.946	0.982	0.991
		Ours	0.055	0.224	2.273	0.091	0.107	0.956	0.984	0.993
	Eigen SfM	S	DORN [15]	0.067	0.295	2.929	0.108	0.130	0.949	0.988
VNL [45]			0.065	0.297	3.172	0.106	0.168	0.945	0.989	0.997
II		DeepV2D [36]	0.050	0.212	2.483	0.089	0.091	0.973	0.992	0.997
		Ours	0.034	0.103	1.919	0.057	0.031	0.989	0.998	0.999

setting of [53], we test our pose estimation accuracy on all 2700 frames of the ‘‘09’’ and ‘‘10’’ sequences, using consecutive frames from the left camera.

MVS, Scenes11, and SUN3D. MVS is collected from several outdoor datasets by [39]. Different from KITTI which is built through video sequences with close scenes, MVS has outdoor scenes from various sources. Scenes11 [39] is a synthetic dataset generated by random shapes and motions. It is therefore annotated with perfect depth and pose, though the images are not realistic. SUN3D [43] provides indoor images with noisy depth and pose annotation. We use the SUN3D dataset post-processed by [39], which discards the samples with a high photo-consistency error.

4.2. Depth Evaluation

We perform depth evaluation on KITTI Depth, MVS, Scenes11, and SUN3D datasets.

KITTI Depth. We compare our framework with both types of deep SfM methods using seven commonly used depth metrics [10]. We also leverage one disparity metric D1-all³ as it measures the precision of the depth estimation. Since the Type **I** methods are self-supervised and they all perform single frame depth estimation in inference, we report the results of the state-of-the-art supervised single image depth estimation methods [15, 45] as they can be viewed as the upper bounds of Type **I** methods.

Quantitative results are shown in Table 1. Although only using a flow estimator trained on synthetic datasets, our method beats all previous methods with a clear margin on various metrics, *e.g.*, 2.273 versus 2.727 in RMSE. Especially, our method largely outperforms DeepV2D although

DeepV2D used ground truth camera pose and five-frame sequences for training. Note that there is a number of frames in the Eigen split that do not strictly satisfy the rigid SfM assumption such as stationary scene. When only keeping the frames that satisfy SfM assumptions, *i.e.*, on the Eigen SfM split, our method achieves even better accuracy, with 3.1% vs 9.1% in D1-all. Fig. 3 illustrates some qualitative results compared with the state-of-the-art supervised single image method [15] and deep SfM method [36].

MVS, Scenes11, and SUN3D. We compare our framework to state-of-the-art Type **II** methods under two-view SfM setting using metrics by [39]. We use the same strategy of iterative depth refinement as [42] in inference for a fair comparison. As shown in Table 2, our method achieves superior performance on all metrics among all three datasets comparing with the previous state-of-the-art Type **II** methods. Fig. 4 provides some qualitative results.

4.3. Camera Pose Estimation

We compare the camera pose estimation accuracy with Type **I** and Type **II** SfM methods on the KITTI VO, MVS, Scenes11, and SUN3D datasets.

KITTI VO. We measure the pose estimation accuracy on relative translational error t_{err} and relative rotational error r_{err} as in [54]. For all results, we align the predicted trajectories to the ground truth via least square optimization [38]. Our method achieves the best pose estimation accuracy with a clear margin compared with the Type **I** SfM methods [53, 2, 32], and full-sequence visual odometry approach [54]. In Fig. 5 we visualize the full sequence odometry trajectories on 9th and 10th sequences. Our results are more aligned with the ground truth trajectories. It is worth noting that our model are only trained on sythetic datasets while the other methods are fine-tuned on the KITTI VO

³Percentage of stereo disparity outliers. We convert the estimated depth to disparities using the focal length and baseline provided by KITTI.

Table 2. **Depth and Pose Estimation Results on MVS, Scenes11, and SUN3D Datasets.** Base-SIFT and Base-Matlab come from [39].

Method	MVS Dataset					Scenes11 Dataset					Sun3D Dataset				
	Depth			Pose		Depth			Pose		Depth			Pose	
	L1-inv	Sc-inv	L1-rel	Rot	Tran	L1-inv	Sc-inv	L1-rel	Rot	Tran	L1-inv	Sc-inv	L1-rel	Rot	Tran
Base-SIFT	0.056	0.309	0.361	21.180	60.516	0.051	0.900	1.027	6.179	56.650	0.029	0.290	0.286	7.702	41.825
Base-Matlab	-	-	-	10.843	32.736	-	-	-	0.917	14.639	-	-	-	5.920	32.298
COLMAP [33]	-	-	0.384	7.961	23.469	-	-	0.625	4.834	10.682	-	-	0.623	4.235	15.956
DeMoN [39]	0.047	0.202	0.305	5.156	14.447	0.019	0.315	0.248	0.809	8.918	0.019	0.114	0.172	1.801	18.811
LS-Net [7]	0.051	0.221	0.311	4.653	11.221	0.010	0.410	0.210	4.653	8.210	0.015	0.189	0.650	1.521	14.347
BANet [35]	0.030	0.150	0.080	3.499	11.238	0.080	0.210	0.130	3.499	10.370	0.015	0.110	0.060	1.729	13.260
DeepSfM [42]	0.021	0.129	0.079	2.824	9.881	0.007	0.112	0.064	0.403	5.828	0.013	0.093	0.072	1.704	13.107
Ours	0.015	0.102	0.068	2.417	3.878	0.005	0.097	0.058	0.276	2.041	0.010	0.081	0.057	1.391	10.757

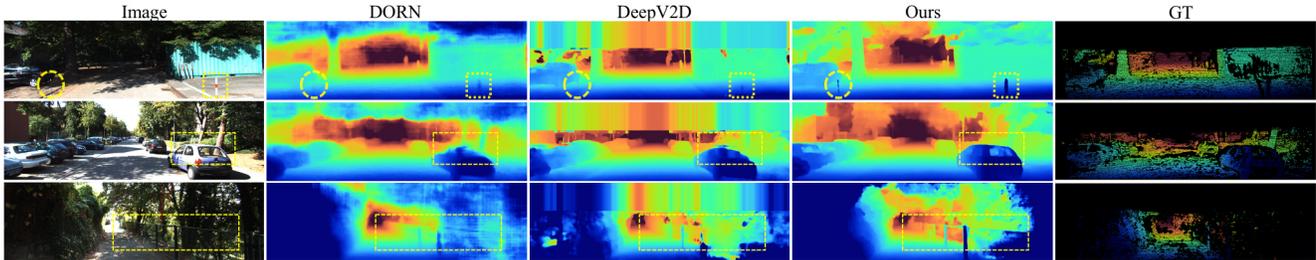


Figure 3. **Qualitative Results on the KITTI Dataset.** The yellow circles and boxes in the top row highlight tiny poles which are captured more accurately by our method.

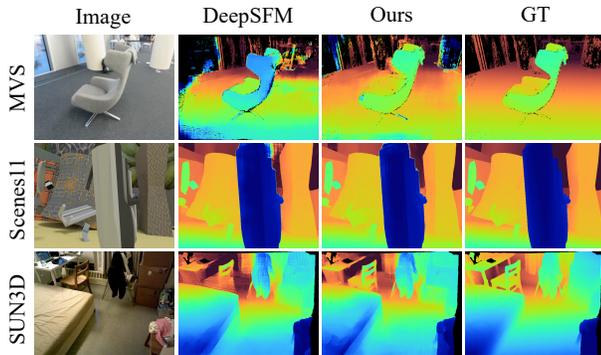


Figure 4. **Qualitative Examples on MVS, Scenes11, and SUN3D Datasets,** where our method consistently achieves better results.

dataset and take more frames to estimate the camera poses.

MVS, Scenes11 and SUN3D. The competitors use ground truth poses to train their pose estimation module on these three datasets, while we use the ground truth poses to fine-tune our optical flow model using Eq. (8). We also report the pose estimation accuracy in Table 2, using the metrics of DeMoN [39]. Our method beats the previous state-of-the-art on all three datasets with a clear margin, *e.g.*, 60.8% better in translation estimation on MVS dataset and 31.5% better in rotation estimation on Scenes11 dataset. Moreover, we verify the effectiveness of rigid flow supervision, Eq (8), in Table 4. With fine-tuning, the translation errors are largely suppressed, and the rotation errors are notably reduced. It is worth noting that our model that was trained on synthetic datasets has already achieved comparable performance with previous methods.

Table 3. **Pose Estimation Accuracy on KITTI VO dataset.** Bold indicates the best. For pose estimation, our method uses an optical flow model trained on synthetic data. The result of GANVO [2] is provided by its author.

Method	Seq. 09		Seq. 10	
	$t_{err}(\%)$	$r_{err}(^{\circ}/100m)$	$t_{err}(\%)$	$r_{err}(^{\circ}/100m)$
SfMLearner [53]	8.28	3.07	12.20	2.96
GANVO [2]	11.52	3.53	11.60	5.17
CCNet [32]	6.92	1.77	7.97	3.11
LTMVO [54]	3.49	1.03	5.81	1.82
Ours	1.70	0.48	1.49	0.55

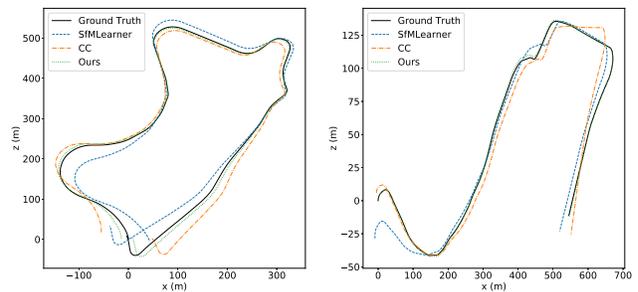


Figure 5. **Visual Trajectory on the KITTI VO dataset.** We compare our method against other deep learning based SfM methods on Seq.09 (Left) and Seq.10 (Right) of KITTI VO dataset.

4.4. Framework Analysis and Justification

Estimating Camera Pose from Optical Flow. There are multiple ways to extract camera pose from optical flow. We consider two kinds of methods: deep regression and the classic five-point algorithm [24] with RANSAC scheme. For deep regression methods, we build a PoseNet similar to the one used in [41] with ResNet50 [19] as the feature back-

Table 4. **The Effect of Optical Flow Fine-tuning.** With the help of rigid flow supervision, our fine-tuned model achieves a much better camera pose result than the model trained on synthetic data.

Model	MVS		Scenes11		SUN3D	
	Rot	Tran	Rot	Tran	Rot	Tran
Our-synthetic	3.637	10.984	0.587	6.617	1.670	12.905
Our-finetune	2.417	3.878	0.276	2.041	1.391	10.757

bone, using image pairs and optical flow as the input. For the five-point algorithm, we use flow matching pairs as the input. We also set a baseline by using SIFT matches. To filter out error matches and outliers, we compare different masking strategies, such as flow uncertainty maps (output of per-pixel softmax operation), learned confidence maps, and SIFT feature locations.

We evaluate these methods on the MVS dataset, see Table 5. Deep regression methods have almost constant performance regardless of different inputs and masking strategies. The best option is to use flow matches with masks based on SIFT feature locations.⁴

Dealing With Misaligned Scales. It is impossible to perfectly recover absolute scales from two-view images. This scale-ambiguity problem will cause trouble if we would like to directly use ground truth depth for supervision or through the widely used scale-invariant loss [10, 39]. We verify the effect of the proposed scale-invariant depth estimation module on the KITTI depth Eigen split. The baseline follows our pipeline but without scale-invariant depth module. It simply uses ℓ_{huber} loss on the estimated depth and the ground depth regardless their scales, which forces the network to implicitly learn a scale. As shown in Table 6, our scale-invariant depth module achieves a very similar accuracy to ‘Oracle Pose’, which is the upper bound of our method. On the other hand, the performance of the scale-invariant loss is similar to the baseline method, which indicates that this loss cannot handle the scale problem.

Scale-invariant Matching on Other Frameworks. The scale-invariant matching is specifically designed for our pipeline to handle the scale ambiguity in depth estimation. Previous deep SfM methods like DeepV2D do not suffer from this problem as they force networks to regress camera poses with scales and then make the depth scaled. That means, these methods cannot benefit from the scale-invariant matching. As a proof, we apply our scale-invariant matching to DeepV2D and test it on the KITTI Eigen dataset. The performance gain is minor: Abs Rel from 0.064 to 0.063 and RMSE from 2.946 to 2.938. Our superior performance benefits from the whole newly proposed deep SfM pipeline rather than a single component. Since all components are tightly coupled in our pipeline, replacing either of them will result in a severe performance drop.

⁴Note that we use SIFT feature *detection* to obtain state-of-the-art results, whereas SIFT feature *matching* performs poorly and is not used.

Table 5. **Estimating Camera Pose from Optical Flow.** We compare different methods to estimate camera pose from optical flow on the MVS dataset. ‘CNN’ represents the pose regression network based on convolutional neural networks with ground truth pose supervision. ‘5-point’ represents the five-point algorithm with RANSAC scheme. We also compare different flow masking strategies here.

Method	Input	Sparse Mask	Rot	Tran
CNN	Color	-	6.652	17.834
CNN	Color + Flow	-	6.437	17.216
CNN	Color + Flow	Uncertainty	6.528	17.107
CNN	Color + Flow	Confidence	6.532	17.511
CNN	Color + Flow	SIFT loc	6.512	17.231
5-point	SIFT matches	-	10.622	29.731
5-point	Flow matches	-	15.673	37.292
5-point	Flow matches	Uncertainty	4.923	12.127
5-point	Flow matches	Confidence	4.614	11.022
5-point	Flow matches	SIFT Loc	2.417	3.878

Table 6. **Dealing with Misaligned Scales.** We compare different strategies to handle the misaligned scales between the estimated depth and ground truth depth on the KITTI Eigen split. ‘Scale Inv Matching’ indicates the scale invariant matching for plane sweeping, ‘Scale Inv Loss’ represents the scale invariant depth loss. The ‘Oracle’ means using the ground truth for both training and inference. Using ground truth pose for training achieves a worse result than the baseline, which verifies the scaling problem.

Strategy	Abs Rel	Sq Rel	RMSE	RMSE _{log}
Baseline	0.089	0.318	3.120	0.129
GT Pose Training	0.121	0.438	3.421	0.175
Scale Inv Loss	0.084	0.302	2.981	0.116
Scale Inv Matching	0.055	0.224	2.273	0.091
Oracle Scale	0.053	0.216	2.271	0.089
Oracle Pose	0.052	0.212	2.269	0.088

5. Conclusion

In this paper, we have revisited the problem of deep neural network based two-view SfM. First, we argued that existing deep learning-based SfM approaches formulate depth estimation or pose estimation as ill-posed problems. Then we proposed a new deep two-view SfM framework that follows the classic well-posed SfM pipeline. Extensive experiments show that our proposed method outperforms all state-of-the-art methods in both pose and depth estimation with a clear margin. In the future, we plan to extend our framework to other SfM problems such as three-view SfM and multi-view SfM, where the loop consistency and temporal consistency could further constrain these already well-posed problems.

Acknowledgements Yuchao Dai was supported in part by National Natural Science Foundation of China (61871325) and National Key Research and Development Program of China (2018AAA0102803). Hongdong Li was supported in part by ACRV (CE140100016), ARC-Discovery (DP 190102261), and ARC-LIEF (190100080) grants. We would like to thank Shihao Jiang, Dylan Campbell, Charles Loop for helpful discussions and Ke Chen for providing field test images from NVIDIA AV cars.

References

- [1] Sameer Agarwal, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M Seitz, and Richard Szeliski. Building rome in a day. *Communications of the ACM*, 54(10):105–112, 2011. 3
- [2] Yasin Almalioglu, Muhamad Risqi U Saputra, Pedro PB de Gusmao, Andrew Markham, and Niki Trigoni. Ganvo: Unsupervised deep monocular visual odometry and depth estimation with generative adversarial networks. In *2019 International conference on robotics and automation (ICRA)*, pages 5474–5480. IEEE, 2019. 3, 6, 7
- [3] M. Bertero, T. Poggio, and V. Torre. Ill-posed problems in early vision. *Proceedings of the IEEE*, 76(8):869–889, Aug. 1988. 1
- [4] Daniel J Butler, Jonas Wulff, Garrett B Stanley, and Michael J Black. A naturalistic open source movie for optical flow evaluation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 611–625. Springer, 2012. 4
- [5] Yuhua Chen, Cordelia Schmid, and Cristian Sminchisescu. Self-supervised learning with geometric constraints in monocular video: Connecting flow, depth, and camera. In *Int. Conf. Comput. Vis.*, pages 7063–7072, 2019. 3, 6
- [6] Xuelian Cheng, Yiran Zhong, Mehrtash Harandi, Yuchao Dai, Xiaojun Chang, Hongdong Li, Tom Drummond, and Zongyuan Ge. Hierarchical neural architecture search for deep stereo matching. *Advances in Neural Information Processing Systems*, 33, 2020. 3
- [7] Ronald Clark, Michael Bloesch, Jan Czarnowski, Stefan Leutenegger, and Andrew J Davison. Ls-net: Learning to solve nonlinear least squares for monocular stereo. *arXiv preprint arXiv:1809.02966*, 2018. 3, 7
- [8] A. J. Davison, I. D. Reid, N. D. Molton, and O. Stasse. MonoSLAM: Real-time single camera SLAM. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 29(6):1052–1067, 2007. 1
- [9] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *Int. Conf. Comput. Vis.*, pages 2758–2766, 2015. 4
- [10] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in neural information processing systems*, pages 2366–2374, 2014. 4, 5, 6, 8
- [11] J. Engel, V. Koltun, and D. Cremers. Direct sparse odometry. In *arXiv:1607.02565*, 2016. 1
- [12] J. Engel, T. Schöps, and D. Cremers. LSD-SLAM: Large-scale direct monocular SLAM. In *European Conference on Computer Vision (ECCV)*, 2014. 1
- [13] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. 1, 3, 4
- [14] Christian Forster, Matia Pizzoli, and Davide Scaramuzza. SVO: Fast semi-direct monocular visual odometry. In *ICRA*, 2014. 1
- [15] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2002–2011, 2018. 6
- [16] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361. IEEE, 2012. 4, 5
- [17] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 1, 3, 4
- [18] T. Hassner and R. Basri. Example based 3D reconstruction from single 2D images. In *Conference on Computer Vision and Pattern Recognition Workshop (CVPRW)*, 2006. 1
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 7
- [20] Berthold KP Horn and Brian G Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981. 4
- [21] Sunghoon Im, Hae-Gon Jeon, Stephen Lin, and In So Kweon. Dpsnet: End-to-end deep plane sweep stereo. *International Conference on Learning Representations*, 2019. 4, 5
- [22] Georg Klein and David Murray. Parallel tracking and mapping for small AR workspaces. In *ISMAR*, 2007. 1
- [23] Dongxu Li, Chenchen Xu, Kaihao Zhang, Xin Yu, Yiran Zhong, Wenqi Ren, Hanna Suominen, and Hongdong Li. Arvo: Learning all-range volumetric correspondence for video deblurring. *arXiv preprint arXiv:2103.04260*, 2021. 3
- [24] Hongdong Li and Richard Hartley. Five-point motion estimation made easy. In *18th International Conference on Pattern Recognition (ICPR’06)*, volume 1, pages 630–633. IEEE, 2006. 1, 4, 7
- [25] Kejie Li, Ravi Garg, Ming Cai, and Ian Reid. Single-view object shape reconstruction using deep shape prior and silhouette. In *BMVC*, 2019. 1
- [26] H Christopher Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293(5828):133–135, 1981. 2, 4
- [27] David G Lowe. Object recognition from local scale-invariant features. In *Proceedings of the seventh IEEE international conference on computer vision*, volume 2, pages 1150–1157. Ieee, 1999. 4
- [28] Reza Mahjourian, Martin Wicke, and Anelia Angelova. Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints. In *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2018. 3
- [29] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4040–4048, 2016. 4
- [30] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. ORB-SLAM: A versatile and accurate monocular SLAM system. *IEEE Transactions on Robotics*, 31(5):1147–1163, 2015. 1, 3

- [31] David Nistér. An efficient solution to the five-point relative pose problem. *IEEE transactions on pattern analysis and machine intelligence*, 26(6):756–770, 2004. 4
- [32] Anurag Ranjan, Varun Jampani, Lukas Balles, Kihwan Kim, Deqing Sun, Jonas Wulff, and Michael J. Black. Competitive collaboration: Joint unsupervised learning of depth, camera motion, optical flow and motion segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2019. 1, 3, 6, 7
- [33] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4104–4113, 2016. 3, 7
- [34] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8934–8943, 2018. 3
- [35] Chengzhou Tang and Ping Tan. Ba-net: Dense bundle adjustment network. *International Conference on Learning Representations*, 2018. 1, 3, 5, 6, 7
- [36] Zachary Teed and Jia Deng. Deepv2d: Video to depth with differentiable structure from motion. *International Conference on Learning Representations*, 2020. 1, 3, 4, 5, 6
- [37] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. *arXiv preprint arXiv:2003.12039*, 2020. 3
- [38] Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns. In *IEEE Trans. Pattern Anal. Mach. Intell.*, pages 376–380. IEEE, 1991. 6
- [39] Benjamin Ummenhofer, Huizhong Zhou, Jonas Uhrig, Nikolaus Mayer, Eddy Ilg, Alexey Dosovitskiy, and Thomas Brox. Demon: Depth and motion network for learning monocular stereo. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5038–5047, 2017. 1, 3, 4, 6, 7, 8
- [40] Jianyuan Wang, Yiran Zhong, Yuchao Dai, Kaihao Zhang, Pan Ji, and Hongdong Li. Displacement-invariant matching cost learning for accurate optical flow estimation. In *Adv. Neural Inform. Process. Syst.*, 2020. 1, 3, 4, 5
- [41] Kaixuan Wang and Shaojie Shen. Flow-motion and depth network for monocular stereo and beyond. *IEEE Robotics and Automation Letters*, 5(2):3307–3314, 2020. 3, 7
- [42] Xingkui Wei, Yinda Zhang, Zhuwen Li, Yanwei Fu, and Xiangyang Xue. Deepsfm: Structure from motion via deep bundle adjustment. *Eur. Conf. Comput. Vis.*, 2020. 3, 4, 6, 7
- [43] Jianxiong Xiao, Andrew Owens, and Antonio Torralba. Sun3d: A database of big spaces reconstructed using sfm and object labels. In *Proceedings of the IEEE international conference on computer vision*, pages 1625–1632, 2013. 6
- [44] Haozhe Xie, Hongxun Yao, Xiaoshuai Sun, Shangchen Zhou, and Shengping Zhang. Pix2Vox: Context-aware 3D reconstruction from single and multi-view images. In *ICCV*, 2019. 1
- [45] Wei Yin, Yifan Liu, Chunhua Shen, and Youliang Yan. Enforcing geometric constraints of virtual normal for depth prediction. In *Int. Conf. Comput. Vis.*, pages 5684–5693, 2019. 6
- [46] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018. 1, 3, 6
- [47] Yiran Zhong, Yuchao Dai, and Hongdong Li. Self-supervised learning for stereo matching with self-improving ability. *arXiv preprint arXiv:1709.00930*, 2017. 1
- [48] Yiran Zhong, Yuchao Dai, and Hongdong Li. Stereo computation for a single mixture image. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, pages 441–456, Cham, 2018. Springer International Publishing. 1
- [49] Yiran Zhong, Yuchao Dai, and Hongdong Li. Efficient depth completion using learned bases. *arXiv preprint arXiv:2012.01110*, 2020. 3
- [50] Yiran Zhong, Pan Ji, Jianyuan Wang, Yuchao Dai, and Hongdong Li. Unsupervised deep epipolar flow for stationary or dynamic scenes. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 12095–12104, 2019. 1
- [51] Yiran Zhong, Hongdong Li, and Yuchao Dai. Open-world stereo video matching with deep rnn. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 3
- [52] Yiran Zhong, Charles Loop, Wonmin Byeon, Stan Birchfield, Yuchao Dai, Kaihao Zhang, Alexey Kamenev, Thomas Breuel, Hongdong Li, and Jan Kautz. Displacement-invariant cost computation for efficient stereo matching. *arXiv preprint arXiv:2012.00899*, 2020. 1
- [53] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1851–1858, 2017. 1, 3, 6, 7
- [54] Yuliang Zou, Pan Ji, , Quoc-Huy Tran, Jia-Bin Huang, and Manmohan Chandraker. Learning monocular visual odometry via self-supervised long-term modeling. In *European Conference on Computer Vision*, 2020. 6, 7